arXiv:1612.00277v1 [cs.PL] 1 Dec 2016

Sparsity Preserving Algorithms for Octagons

Jacques-Henri Jourdan

MPI-SWS, Inria Paris

Abstract

Known algorithms for manipulating octagons do not preserve their sparsity, leading typically to quadratic or cubic time and space complexities even if no relation among variables is known when they are all bounded. In this paper, we present new algorithms, which use and return octagons represented as weakly closed difference bound matrices, preserve the sparsity of their input and have better performance in the case their inputs are sparse. We prove that these algorithms are as precise as the known ones.

1 Introduction

In order to capture numerical properties of programs, static analyzers use *numerical abstract domains*. The choice of a numerical abstract domain in a static analyzer is a compromise between *precision*, the ability of capturing complex numerical properties, and performance. Non-relational abstract domains, such as intervals [6], are very efficient but relatively imprecise: they cannot represent relations between program variables. On the other hand, in order to capture numerical relations between program variables, one can express them as linear inequalities. This class of relational numerical abstract domain is composed of *linear abstract domains*. A linear abstract domain corresponds to a different precision vs. performance trade-off: they range from the less precise, efficient ones such as zones [13], pentagons [12] or octagons [13,14] to the more precise, costly ones, such as subpolyhedra [11], octahedra [5], two variables per inequalities [16], zonotopes [15] or general polyhedra [8].

In particular, the Octagon abstract domain [13,14] accurately represents many of the variable relationships appearing in a program, while being still reasonably fast (all the operations have quadratic or cubic complexity on the number of variables). It is very popular in the static analysis community, which explains why algorithmic improvements [3,1,17] and precision improving variants [4] are regularly published.

As reported by the designers of Astrée [7], its quadratic or cubic performances still make it unusable as-is with a reasonable number of variables. Indeed, the

[1] This work was supported by Agence Nationale de la Recherche, grant ANR-11-INSE-003.

This paper is electronically published in
Electronic Notes in Theoretical Computer Science
URL: `www.elsevier.com/locate/entcs`

data structures typically used to represent octagonal abstract values, i.e., strongly closed difference bound matrices, have a quadratic size in the number of variables for which an upper or lower bound is known. A common solution is the use of *variable packing* [13, §8.4.2], where the Octagon abstract domain is only used on small packs of variables. The downside of packing is that no relation is stored between variables that are not in the same pack. A variant of packing has been introduced to mitigate the imprecision [2], but loss in precision can still occur.

The problem of the performance of octagons has already been studied: in particular, Singh et al. [17] proposed an implementation of the Octagon abstract domain optimized in the case its representation is sparse. But they do not address the fact that it is dense as soon as interval bounds are known for many variables, and we anticipate that, for this reason, the sparsity is very low in their implementation.

Instead, in this paper, we propose to use new algorithms for the Octagon abstract domain: these algorithms work on a sparse representation for octagons, so that the cost of the analysis of two independent sets of variables is the sum of the costs of the analyses of the two sets of variables, taken independently. Our algorithms have the same precision as the traditional ones. Our main idea is the following: in order to ensure an optimal precision of all the operations, the data structures representing octagons, difference bound matrices, are usually kept *strongly closed*: that is, algorithms make sure that any returned difference bound matrix is a best abstraction. However, most often, strongly closed difference bound matrices are dense because of the necessary *strengthening* step. In this paper, we propose to weaken the maintained invariant on difference bound matrices and to keep them *weakly closed* hence skipping the strengthening step. Weakly closed difference bound matrices are not necessarily dense, so that we can use sparse data structures to represent them. We prove that some algorithms can be kept unchanged to work on weakly closed difference bound matrices without losing any precision and give new algorithms for the other operations.

We begin by preliminary definitions in §2. In §3, we describe and prove the soundness and relative precision of our new algorithms. We conclude in §4.

2 Definitions

Let $\mathbb{V}_+$ be a finite set of variables. We call a *regular environment* a function from $\mathbb{V}_+$ to $\mathbb{R}$. A regular environment represents the numerical state of a program. The role of the Octagon abstract domain is to approximate sets of regular environments ρ. To that end, the abstract domain of octagons stores a set of inequalities of the following form:

$$\pm\rho(u) \pm \rho(v) \leq Cst_{uv} \qquad u, v \in \mathbb{V}_+ \tag{1}$$

This corresponds to giving bounds to sums and differences of values of ρ. Moreover, if we use twice the same variable with the same sign, we see that, using such constraints, we can express interval constraints over values of an environment [13].

In order to handle in a unified way all the different combinations of signs in these constraints, we introduce the set $\mathbb{V}_\pm$ of *signed variables*. Signed variables are of two kinds: they are either usual variables from $\mathbb{V}_+$, called *positive variables* in the context of signed variables, or their opposites form, *negative variables*. We equip

$\mathbb{V}_\pm$ with an involutive operator, associating to each signed variable v its opposite $\overline{v}$, such that v is positive if and only if $\overline{v}$ is negative.

Regular environments are canonically extended to signed variables by taking $\rho(\overline{u}) = -\rho(u)$. More generally, we define *irregular environments* as functions σ from $\mathbb{V}_\pm$ to $\mathbb{R}$, and consider the set of regular environments as a subset of the set of irregular environments. Regular environments are exactly irregular environments ρ that satisfy the property $\forall v, \rho(\overline{v}) = -\rho(v)$.

Using this new formalism, octagonal constrains of the form (1) can be seen as upper bounds on differences of values of ρ, a regular environment:

$$\rho(u) - \rho(v) \leq Cst_{uv} \qquad u, v \in \mathbb{V}_\pm \tag{2}$$

This has two benefits: first, all the different kinds of constraints allowed by (1) get factored out as one simpler form. Second, we can see these constraints as constraints on irregular environments, and further constrain them as being regular: we see that the study of the Octagon abstract domain starts by the study of a simpler abstract domain, where only differences of variables are bounded. The set of constraints, called *potential constraints*, of such an abstract domain is well studied in the linear optimization literature, because it corresponds to the well-known shortest path problem in a weighted directed graph.

Such a set of constraints is represented as difference bound matrices: a *difference bound matrix*, or *DBM*, is a matrix $(B_{uv})_{(u,v)\in\mathbb{V}_\pm^2}$ of elements of $\mathbb{R} \cup \{+\infty\}$. The meaning of these constraints is given by two concretization functions γ_{pot} and γ_{oct}, that associate to a DBM the set of irregular or regular environments, respectively, satisfying all the constraints:

$$\gamma_{pot}(B) = \{\sigma : \mathbb{V}_\pm \to \mathbb{R} \mid \forall uv \in \mathbb{V}_\pm,\ \sigma(u) - \sigma(v) \leq B_{uv}\} \tag{3}$$
$$\gamma_{oct}(B) = \{\rho \in \gamma_{pot}(B) \mid \forall u \in \mathbb{V}_\pm,\ \rho(\overline{u}) = -\rho(u)\} \tag{4}$$

Example 2.1 Consider $\mathbb{V}_+ = \{x; y; z\}$ a set of three (positive) variables. The set of signed variables is $\mathbb{V}_\pm = \{x; \overline{x}; y; \overline{y}; z; \overline{z}\}$. Let A be the DBM such that $A_{x\overline{x}} = 1$, $A_{\overline{y}y} = 3$, $A_{y\overline{z}} = 1$ and $A_{uv} = +\infty$ for all the other entries. The set $\gamma_{oct}(A)$ contains all the environments $\rho : \mathbb{V}_\pm \to \mathbb{R}$ such that:

- $\forall u \in \mathbb{V}_\pm,\ \rho(\overline{u}) = -\rho(u)$
- $\rho(x) \leq 1/2$, $-\rho(y) \leq 3/2$ and $\rho(y) + \rho(z) \leq 1$

This concretization is assimilated to the set of environments $\rho : \mathbb{V}_+ \to \mathbb{R}$ over positive variables such that $\rho(x) \leq 1/2$, $-\rho(y) \leq 3/2$ and $\rho(y) + \rho(z) \leq 1$.

We denote as $\overset{\sharp}{\leq}$ the natural order relation over DBMs, defined as follows:

$$A \overset{\sharp}{\leq} B \Leftrightarrow \forall uv \in \mathbb{V}_\pm,\ A_{uv} \leq B_{uv} \tag{5}$$

The following easy lemma states that this order relation makes γ_{pot} and γ_{oct} increasing, which makes $\overset{\sharp}{\leq}$ a good candidate for a comparison operator of the Octagon abstract domain [2]:

[2] As we see in §3.1, this is *not* the order relation we use as a comparison operator in our implementation of the Octagon abstract domain.

Lemma 2.2 *Let A and B be two DBMs such that $A \overset{\sharp}{\leq} B$. Then, we have:*

$$\gamma_{pot}(A) \subseteq \gamma_{pot}(B) \qquad\qquad \gamma_{oct}(A) \subseteq \gamma_{oct}(B)$$

For any non-empty set S of irregular environments, there exists a minimal (in the sense of $\overset{\sharp}{\leq}$) DBM that approximates it. That is, there exists a minimal DBM $\alpha(S)$ such that $S \subseteq \gamma_{pot}(\alpha(S))$. This property follows immediately from the definition of α:

$$\alpha(S)_{uv} = \sup_{\sigma \in S} \{\sigma(u) - \sigma(v)\} \tag{6}$$

This function α is called the abstraction function. We can easily see that α is an increasing function. Moreover, α does not only return best abstractions for γ_{pot}, but also for γ_{oct}: if the set S contains only regular environments, we can see that $\alpha(S)$ is also the minimal DBM such that $S \subseteq \gamma_{oct}(\alpha(S))$. In fact, it is easy to see that $\overset{\sharp}{\leq}$ defines a complete lattice over DBMs extended with a bottom element, and that the pairs (α, γ_{pot}) and (α, γ_{oct}) form Galois connections.

2.1 Closure and Strong Closure

Many DBMs have the same concretization. This is a problem, because the abstract environments that we manipulate are therefore not necessarily the most precise ones, and this can lead to imprecision. Thus, usually, an implementation of the Octagon abstract domain maintains the invariant that it only manipulates "canonical" forms of DBMs, such that $B = \alpha(\gamma_{oct}(B))$. Such "canonical" DBMs are always the best possible representative over all the DBMs with the same concretization.

An important fact is that we can characterize best abstractions using the values they contain, and that we have algorithms to compute them. We expose these characterizations, together with these algorithms. Moreover, we give a weaker closedness condition over DBMs, that does not ensure canonicity, but that allows better algorithms without loss of precision.

2.1.1 Best abstractions for γ_{pot}

A first step is to remark that canonical DBMs always have null diagonal values. Moreover, canonical DBMs should always verify the triangular inequality. We call such DBMs *closed* DBMs:

Definition 2.3 [Closed DBM] A *closed* DBM is a DBM B verifying the two following properties:

- $\forall v \in \mathbb{V}_{\pm},\ B_{vv} = 0$
- $\forall uvw \in \mathbb{V}_{\pm},\ B_{uw} \leq B_{uv} + B_{vw}$

Closed DBMs are exactly best abstractions for γ_{pot} [13, Theorem 3.3.6]. Hence, closed DBMs always have non-empty concretizations. We do not detail here the algorithm used to detect the emptiness of the concretization of a DBM and to compute closures: instead, we refer the interested reader to previous work [13,1].

Example 2.4 The closure $\alpha(\gamma_{pot}(A))$ of the DBM A as defined in Example 2.1 contains the following additional finite entries:

- $\forall u \in \mathbb{V}_{\pm},\ \alpha(\gamma_{pot}(A))_{uu} = 0$
- $\alpha(\gamma_{pot}(A))_{\overline{y}\,\overline{z}} = 4$ (corresponding to the constraint $\rho(z) - \rho(y) \leq 4$).

2.1.2 Best abstractions for γ_{oct}

We now refine the notion of closure to canonical forms for γ_{oct}. It is easy to see that, for any non-empty set S of regular environments, $\alpha(S)_{uv} = \alpha(S)_{\overline{v}\,\overline{u}}$. Thus, canonical DBMs for γ_{oct} will verify the *coherence* property:

Definition 2.5 [Coherent DBM] A DBM B is *coherent* when:

$$\forall uv \in \mathbb{V}_{\pm},\ B_{uv} = B_{\overline{v}\,\overline{u}}$$

Moreover, matrix elements of the form $B_{u\overline{u}}$ (for $u \in \mathbb{V}_{\pm}$) impose interval constraints on values of ρ. These interval constraints can be combined to entail constraints on any difference of values of ρ. For this reason, canonical forms for γ_{oct} will verify the following strong closedness property:

Definition 2.6 [Strongly closed DBM] A DBM B is *strongly closed* when it is closed and coherent and:

$$\forall uv \in \mathbb{V}_{\pm},\ B_{uv} \leq \frac{B_{u\overline{u}} + B_{\overline{v}v}}{2}$$

This condition is necessary and sufficient: strong closedness characterizes canonical DBMs for γ_{oct}.

Theorem 2.7 *Let B be a DBM. The two following properties are equivalent:*

(i) B is strongly closed

(ii) $\gamma_{oct}(B) \neq \emptyset$ and $B = \alpha(\gamma_{oct}(B))$

Proof. See, e.g., [13, Theorems 4.3.2 and 4.3.3]. □

Usually [1], to compute strong closure, one first ensures that the given matrix is coherent, then computes a closure (i.e., a canonical representative in the sense of γ_{pot}), and, finally, performing a so-called *strengthening step* [3]:

Definition 2.8 [Strengthening] Let B be a DBM. The *strengthening* of B, noted $^\sharp\mathcal{S}(B)$ is defined by:

$$^\sharp\mathcal{S}(B)_{uv} = \min\left\{\frac{B_{u\overline{u}} + B_{\overline{v}v}}{2}; B_{uv}\right\}$$

The following theorem states the correctness of the strong closure algorithm sketched above, consisting in computing a closure followed by a strengthening:

Theorem 2.9 *Let B be a coherent DBM with $\gamma_{oct}(B) \neq \emptyset$. Then:*

$$\alpha(\gamma_{oct}(B)) = {}^\sharp\mathcal{S}(\alpha(\gamma_{pot}(B)))$$

In particular, if B is coherent and closed, then $^\sharp\mathcal{S}(B)$ is strongly closed.

Proof. See, e.g., [10, Theorem 8.2.7]. □

[3] This is actually an improvement of the method described initially by Miné [13].

Example 2.10 In order to consider the strong closure of the DBM A as defined in Example 2.1, we first need to make it coherent: let $\tilde{A}$ be the DBM containing the same entries as A, except that $\tilde{A}_{z\overline{y}} = 1$.

The closure of $\tilde{A}$ contains the following addional finite entries:

- $\forall u \in \mathbb{V}_{\pm},\ \alpha(\gamma_{pot}(\tilde{A}))_{uu} = 0$
- $\alpha(\gamma_{pot}(\tilde{A}))_{zy} = \alpha(\gamma_{pot}(A))_{\overline{y}\,\overline{z}} = 4$ (corresponding to the constraint $\rho(z) - \rho(y) \leq 4$)
- $\alpha(\gamma_{pot}(\tilde{A}))_{z\overline{z}} = 5$ (corresponding to the constraint $\rho(z) \leq 5/2$).

The strong closure $\alpha(\gamma_{oct}(\tilde{A}))$ is then obtained by strengthening $\alpha(\gamma_{pot}(\tilde{A}))$. The strengthening operation creates the following new entries:

- $\alpha(\gamma_{oct}(\tilde{A}))_{xy} = \alpha(\gamma_{oct}(\tilde{A}))_{\overline{y}\,\overline{x}} = 2$
- $\alpha(\gamma_{oct}(\tilde{A}))_{x\overline{z}} = \alpha(\gamma_{oct}(\tilde{A}))_{z\overline{x}} = 3$.

2.2 Weak Closedness

Usually, the implementations of the Octagon abstract domain maintain all DBMs strongly closed, so that maximal information is known when performing an abstract operation. However, this breaks sparsity: indeed, matrix elements of the form $B_{u\overline{u}}$ are non-relational interval bounds on the variables: as we expect many variables to be bounded, the strengthening step gives finite bounds for many DBM cells, and a strengthened DBM loses most of the sparsity. In general, a DBM has a quadratic size in the number of variables, and therefore this loss of sparsity is costly. Previous attempts at improving performances using sparsity [17] did not make this observation. We believe that, when using these implementations, DBMs quickly become dense, hence reducing the efficiency of sparse algorithms.

In our algorithms, we propose to skip the strengthening step: instead of maintaining the invariant that all the manipulated DBMs are strongly closed, we maintain the invariant that they are *weakly* closed:

Definition 2.11 [Weakly closed DBM] Let B be a DBM. We say that B is *weakly closed* when any of the two following *equivalent* statements hold:

(i) B has a null diagonal and $^\sharp\mathcal{S}(B)$ is strongly closed;

(ii) B has a null diagonal, $^\sharp\mathcal{S}(B)$ is coherent, and:

$$\forall uvw,\ {}^\sharp\mathcal{S}(B)_{uw} \leq B_{uv} + B_{vw} \tag{7}$$

Proof. The proof of equivalence of the definitions is in [10, Definition 8.2.5]. □

In order to make sure we do not lose precision, we will prove for each of those operators that it computes abstract values with the same concretization as with the usual algorithms. Equivalently, we prove that the strengthening of the abstract values computed by our operators are equal to the abstract values computed by the usual operators on the strengthened parameters.

A weakly closed DBM is neither necessarily strongly closed nor closed. However, a closed and coherent DBM is always weakly closed: this helps us easily building weakly closed DBMs from arbitrary sets of octagonal constraints.

Example 2.12 Continuing on the definitions of Example 2.10, $\alpha(\gamma_{pot}(\tilde{A}))$ is closed and coherent hence weakly closed. This DBM contains no entry relating the variable x and the other variables. This is an improvement in sparsity compared to the strong closure $\alpha(\gamma_{oct}(\tilde{A}))$. To the best of our knowledge, this opportunity is not leveraged by previously known algorithms, such as [17].

This notion of weak closedness has been introduced by Bagnara et al. [1, Appendix A] as an intermediate notion for proving the correctness of the tight closure algorithm (see §3.5). To the best of our knowledge, the use of weak closedness as an invariant for manipulating sparse DBMs is an original result of our work.

3 Operations on Difference Bound Matrices

The abstract domain of octagons defines several operations manipulating difference bound matrices. They include lattice operations, like comparison and join, and abstract transfer functions, which model state change in the program.

In this section, we recall the standard definition of these operations, and give the new sparsity-preserving definition on weakly closed DBMs. All these algorithms preserve the sparsity and weak closedness of DBMs and can be proved to be as precise as the standard ones. More precisely, we claim that they always return DBMs whose strengthening equals the DBMs that would have been returned by the traditional algorithms. The implementation of the widening operation, detailed in [10, Section 8.2.7], is more complex and omitted by lack of space.

3.1 Comparison

In order to use octagons in a static analyzer, we need to define a comparison operator, taking two DBMs and returning a Boolean. If this Boolean is `true`, then we have the guarantee that the concretization of the first operand is included in that of the second operand.

A good candidate is ${}^\sharp\leq$, the natural order relation between DBMs. Its soundness is guaranteed by the monotonicity of γ_{oct}. In usual implementations of the Octagon abstract domain, DBMs are kept strongly closed, hence this operator is actually as precise as possible: it returns `true` if and only if the concretizations are included.

However, in the setting of weakly closed DBMs, this property does not hold. In order not to lose precision while still using sparse DBM, we need another comparison operator that strengthens the bounds of the left operand when they do not entail the right operand:

Definition 3.1 [Weakly closed comparison] Let A and B be DBMs. The weakly closed comparison of A and B, noted $A \,{}^\sharp\leq_{weak}\, B$ is defined by:

$$A \,{}^\sharp\leq_{weak}\, B \equiv \bigwedge_{\substack{u,v\in\mathbb{V}_\pm \\ B_{uv}<+\infty}} A_{uv} \leq B_{uv} \;\vee\; \frac{A_{u\overline{u}} + A_{\overline{v}v}}{2} \leq B_{uv}$$

That is, for every finite bound on B, we first check whether it is directly entailed by the corresponding bound in A, and then try to entail it using non-relational

bounds. The following theorem states that it implements the comparison on concretizations, hence we can use it in a sparse context without losing precision:

Theorem 3.2 *Let A a weakly closed DBM and B any DBM. The two following statements are equivalent:*

(i) $\gamma_{oct}(A) \subseteq \gamma_{oct}(B)$

(ii) $A \,{}^\sharp{\leq}_{weak}\, B$

Proof. See, e.g., [10, Theorem 8.2.9]. □

3.2 Forgetting Variables

An important operation provided by the Octagon abstract domain is `forget`. When given a DBM and a variable v, it returns another DBM where all the information on v has been forgotten. Its concrete and abstract definitions are given by:

Definition 3.3 [Concrete forgetting] Let $x \in \mathbb{V}_+$ and S be a set of regular environments. We define:

$$\mathcal{F}^x_{oct}(S) = \{\sigma + [x \Rightarrow r; \overline{x} \Rightarrow -r] \mid \sigma \in S, r \in \mathbb{R}\}$$

Definition 3.4 [Abstract forgetting]

(i) Let $x \in \mathbb{V}_\pm$ and B be a DBM. We define ${}^\sharp\mathcal{F}^x_{pot}(B)$ the DBM such that:

$${}^\sharp\mathcal{F}^x_{pot}(B)_{uv} = \begin{cases} 0 & \text{if } u = v = x \\ +\infty & \text{otherwise if } u = x \text{ or } v = x \\ B_{uv} & \text{otherwise} \end{cases}$$

(ii) Let $x \in \mathbb{V}_+$ and B a DBM. We define:

$${}^\sharp\mathcal{F}^x_{oct}(B) = {}^\sharp\mathcal{F}^x_{pot}({}^\sharp\mathcal{F}^{\overline{x}}_{pot}(B))$$

It is a known result from the Octagon literature [13, Theorems 3.6.1 and 4.4.2] that ${}^\sharp\mathcal{F}^x_{oct}$ is sound when applied to any DBM. Moreover, when applied to any strongly closed DBM, it is exact and returns a strongly closed DBM. To these properties, we add similar properties for weak closedness, that let us use ${}^\sharp\mathcal{F}^x_{oct}$ as-is for weakly closed DBMs without loss of precision:

Theorem 3.5 *Let B be a weakly closed DBM and $x \in \mathbb{V}_+$. We have:*

(i) ${}^\sharp\mathcal{S}({}^\sharp\mathcal{F}^x_{oct}(B)) = {}^\sharp\mathcal{F}^x_{oct}({}^\sharp\mathcal{S}(B))$

(ii) $\mathcal{F}^x_{oct}(\gamma_{oct}(B)) = \gamma_{oct}({}^\sharp\mathcal{F}^x_{oct}(B))$

(iii) ${}^\sharp\mathcal{F}^x_{oct}(B)$ *is weakly closed*

Proof. See [10, Theorem 8.2.11]. □

3.3 Join

The usual join operator on DBMs is the least upper bound operator for ${}^\sharp\leq$:

Definition 3.6 [DBM least upper bound] Let A and B be two DBMs. The least upper bound ${}^\sharp\cup$ on DBMs is defined by:

$$\forall uv,\ (A\ {}^\sharp\cup\ B)_{uv} = \max\{A_{uv}\ ;\ B_{uv}\}$$

The order relation ${}^\sharp\leq$ and the operator ${}^\sharp\cup$ clearly form an upper semi-lattice, thus usual properties on Galois connections hold, providing the usual results on the soundness and precision of this operator: ${}^\sharp\cup$ is sound, and, if given strongly closed DBMs, it returns the best strongly closed DBM approximating the concrete union.

For weakly closed DBMs, even though ${}^\sharp\cup$ is sound, it would possibly lose precision when applied to non-strongly closed DBMs. For example, the weakly closed DBM A represents the two following inequalities on positive variables x and y:

$$x + x \leq 1 \qquad\qquad y + y \leq 0$$

The weakly closed DBM B, in turn, represents the two following inequalities:

$$x + x \leq 0 \qquad\qquad y + y \leq 1$$

The inequality $x + y \leq 1/2$ is not present in A nor in B, even though it is in ${}^\sharp\mathcal{S}(A)$ and in ${}^\sharp\mathcal{S}(B)$. As a result, $A\ {}^\sharp\cup\ B$ contains the inequalities $x + x \leq 1$ and $y + y \leq 1$, but does not entail $x + y \leq 1/2$, which is entailed however by ${}^\sharp\mathcal{S}(A)\ {}^\sharp\cup\ {}^\sharp\mathcal{S}(B)$.

The rationale behind this example is that a join can create some amount of relationality that was not present in one or both operands. Our operator has to reflect this fact. Care should be taken, however, not to break the sparsity of the operands by introducing spurious finite values in the matrix. Our join for weakly closed DBMs is defined as follows:

Definition 3.7 [Weakly closed join for octagons] Let A and B be two weakly closed DBMs. We take, for $u, v \in \mathbb{V}_\pm$, $B^{1/2}_{uv} = \frac{B_{u\overline{u}} + B_{\overline{v}v}}{2}$ and $A^{1/2}_{uv} = \frac{A_{u\overline{u}} + A_{\overline{v}v}}{2}$. The weakly closed join ${}^\sharp\cup_{weak}$ is defined in two steps:

(i) We first define $A\ {}^\sharp\cup^0_{weak}\ B$. Let $u, v \in \mathbb{V}_\pm$. We define:

$$(A\ {}^\sharp\cup^0_{weak}\ B)_{uv} = \begin{cases} A_{uv} & \text{if } A_{uv} = B_{uv} \\ B_{uv} & \text{if } A_{uv} < B_{uv} \leq B^{1/2}_{uv} \\ \max\{A_{uv}\ ;\ B^{1/2}_{uv}\} & \text{if } A_{uv} < B_{uv} \wedge B^{1/2}_{uv} < B_{uv} \\ (B\ {}^\sharp\cup^0_{weak}\ A)_{uv} & \text{if } A_{uv} > B_{uv} \end{cases}$$

(ii) Let $u, v \in \mathbb{V}_\pm$. We define:

$$(A\ {}^\sharp\cup_{weak}\ B)_{uv} = \begin{cases} \min\left\{ \begin{array}{c} (A\ {}^\sharp\cup^0_{weak}\ B)_{uv} \\ \max\left\{A^{1/2}_{uv}\ ;\ B^{1/2}_{uv}\right\} \end{array} \right\} & \text{if } A_{u\overline{u}} < B_{u\overline{u}} \wedge A_{\overline{v}v} > B_{\overline{v}v} \\ & \text{or } A_{u\overline{u}} > B_{u\overline{u}} \wedge A_{\overline{v}v} < B_{\overline{v}v} \\ (A\ {}^\sharp\cup^0_{weak}\ B)_{uv} & \text{otherwise} \end{cases}$$

The first step can be computed by iterating over all the matrix elements that are different in A and B. This first step thus preserves the sparsity, and consumes

computing time only for variables that are different in both branches. The second step can be computed efficiently by first collecting in a list all the variables u for which $A_{u\overline{u}} < B_{u\overline{u}}$ and, in another list, all those for which $B_{u\overline{u}} < A_{u\overline{u}}$. By iterating over the two lists, we can efficiently modify only the cells meeting the given condition. It should be noted that we break in the second step only the sparsity that needs to be broken, as the modified cells correspond to the cases where the join create new relational information (as in the example above).

The following theorem states that this modified join operator can be used on weakly closed DBMs without losing precision or soundness:

Theorem 3.8 *Let A and B be two weakly closed DBMs. We have:*

(i) $\sharp\mathcal{S}(A \mathbin{\sharp\cup_{weak}} B) = \sharp\mathcal{S}(A) \mathbin{\sharp\cup} \sharp\mathcal{S}(B)$

(ii) $\gamma_{oct}(A \mathbin{\sharp\cup_{weak}} B) = \gamma_{oct}(\alpha(\gamma_{oct}(A) \cup \gamma_{oct}(B)))$

(iii) $A \mathbin{\sharp\cup_{weak}} B$ *is weakly closed*

Proof. See [10, Theorem 8.2.13]. □

3.4 Assuming Constraints

An important operation for abstract domains is the `assume` primitive, which refines the internal state of an abstract domain using a new assumption over the set of approximated environments. In this section, we only consider the cases where this operation is exact, i.e., it does not lead to any approximation. These cases amount to assuming that $\rho(x) - \rho(y) \leq C$, for $C \in \mathbb{R}$ and x and y two variables. In order to deal with arbitrary linear inequalities or even arbitrary arithmetical constraints, it is necessary to write some supporting module for the Octagon domain that will translate arbitrary constraints into exact ones. Such a support module is out of the scope of this paper: we refer the reader to [13] for more detail. Moreover, note that the combination of `assume` together with the `forget` let us emulate variable assignment [4], hence we do not detail variable assignment in this paper.

We give the `assume` primitive in two versions: one adapted to γ_{pot}, and one adapted to γ_{oct}. We first give the concrete semantics of this operation, which is the same for irregular and regular environments:

Definition 3.9 [Assuming constraints in the concrete] Let $C \in \mathbb{R}$, $x, y \in \mathbb{V}_{\pm}$ and S be a set of irregular environments. We define:

$$\mathcal{A}^{x-y\leq C}(S) = \{\sigma \in S \mid \sigma(x) - \sigma(y) \leq C\}$$

It is easy to see that we can reflect exactly this operation in DBMs. Indeed, it suffices to change the cell corresponding to the new constraint, if the old value is larger than the new one. However, this does not maintain any kind of closedness, whether it be the normal closure, the strong closure or the weak closedness. As a result, it is necessary to run a closure algorithm when inserting the new constraint. These algorithms are costly (i.e., cubic complexity), and do not leverage the fact that the input matrix is already almost closed. For this reason, *incremental closure*

[4] An efficient implementation would however use a specific, optimized implementation for assignments.

algorithms have been developed, with quadratic complexity. We give here a slightly different presentation of these algorithms as the one originally given by Miné [13]:

Definition 3.10 [Assuming constraints in the abstract] Let $C \in \mathbb{R}$, B be a DBM and $x, y \in \mathbb{V}_\pm$.

(i) We define $\sharp\mathcal{A}_{pot}^{x-y\leq C}(B)$ the DBM such that, for $u, v \in \mathbb{V}_\pm$:

$$\sharp\mathcal{A}_{pot}^{x-y\leq C}(B)_{uv} = \min\{B_{uv} \ ; \ B_{ux} + C + B_{yv}\}$$

(ii) If $x, y \in \mathbb{V}_+$, we define $\sharp\mathcal{A}_{weak}^{x-y\leq C}(B)$ and $\sharp\mathcal{A}_{oct}^{x-y\leq C}(B)$ as:

$$\sharp\mathcal{A}_{weak}^{x-y\leq C}(B) = \sharp\mathcal{A}_{pot}^{\overline{y}-\overline{x}\leq C}(\sharp\mathcal{A}_{pot}^{x-y\leq C}(B))$$
$$\sharp\mathcal{A}_{oct}^{x-y\leq C}(B) = \sharp\mathcal{S}(\sharp\mathcal{A}_{weak}^{x-y\leq C}(B))$$

It is well-known [10, Theorem 8.2.14] that $\sharp\mathcal{A}_{oct}^{x-y\leq C}$ is sound and exact when applied to a DBM with a null diagonal. When applied to a strongly closed DBM B with $0 \leq C + B_{yx}$, the result is strongly closed. Therefore, an implementation of the `assume` primitive in the strongly closed setting first checks whether $0 \leq C + B_{yx}$. If so, it returns $\sharp\mathcal{A}_{oct}^{x-y\leq C}$; otherwise it returns $\bot$.

In particular, when applied to weakly closed DBMs, $\sharp\mathcal{A}_{oct}^{x-y\leq C}$ is sound and exact, since weakly closed DBMs have null diagonals. However, because this operator uses $\sharp\mathcal{S}$, it breaks sparsity. The advantage of using weakly closed DBMs is that, in the setting of weakly closed DBMs, $\sharp\mathcal{S}$ is no longer needed: $\sharp\mathcal{A}_{weak}^{x-y\leq C}$ can be used as-is, provided the implementation additionally checks that $0 \leq 2C + B_{y\overline{y}} + B_{\overline{x}x}$. The following theorem summarizes this result, and justifies the use of this transfer function in the context of sparse DBMs without loss of precision:

Theorem 3.11 *Let $C \in \mathbb{R}$, B a weakly closed DBM and $x, y \in \mathbb{V}_+$. We have:*

(i) *If $0 \leq 2C + B_{y\overline{y}} + B_{\overline{x}x}$, then $\sharp\mathcal{S}(\sharp\mathcal{A}_{weak}^{x-y\leq C}(B)) = \sharp\mathcal{A}_{oct}^{x-y\leq C}(\sharp\mathcal{S}(B))$*

(ii) *$\gamma_{oct}(\sharp\mathcal{A}_{weak}^{x-y\leq C}(B)) = \mathcal{A}^{x-y\leq C}(\gamma_{oct}(B))$*

(iii) *If B is weakly closed, the following statements are equivalent:*
  (i) $\mathcal{A}^{x-y\leq C}(\gamma_{oct}(B)) \neq \emptyset$
  (ii) $0 \leq \sharp\mathcal{A}_{weak}^{x-y\leq C}(B)_{xx}$
  (iii) $0 \leq C + B_{yx}$ and $0 \leq 2C + B_{y\overline{y}} + B_{\overline{x}x}$
  (iv) $\sharp\mathcal{A}_{weak}^{x-y\leq C}(B)$ is weakly closed.

Proof. See, e.g., [10, Theorem 8.2.15]. □

3.5 Tightening

Miné [13] and Bagnara et al. [1] study the case of the Octagon abstract domain when the considered environments take only values in $\mathbb{Z}$: in contrast with the previous sections, in this case, the strongly closed DBMs are not all canonical, so that modified algorithms need to be used. We explain here that the use of the weakly closed setting is compatible with the integer case. To this end, we define a different concretization function, $\gamma_{oct}^{\mathbb{Z}}$, that concretizes to integer environments:

Definition 3.12 [Integer concretization of octagons] Let B be a DBM. We define:

$$\gamma_{oct}^{\mathbb{Z}}(B) = \{\rho \in \gamma_{oct}(B) \mid \forall u \in \mathbb{V}_+,\ \rho(u) \in \mathbb{Z}\}$$

If we consider only integer environments, best abstractions have a slightly stronger characterization. Such DBM are said *tightly closed*. We also define the notion of *weakly tightly closed* DBMs, which is the analog of *weakly closed* DBMs for the integer case:

Definition 3.13 [Tight closure] Let B be a DBM. B is *tightly closed* (respectively weakly tightly closed) when:

- B is strongly closed (respectively weakly closed)
- $\forall uv \in \mathbb{V}_\pm,\ B_{uv} \in \mathbb{Z}$
- $\forall u \in \mathbb{V}_\pm,\ \frac{B_{u\overline{u}}}{2} \in \mathbb{Z}$

Tightly closed DBMs are exactly best abstractions for integer environments [10, Theorem 8.2.17]. Bagnara et al. [1, §6] give efficient algorithms for computing the tight closure of a DBM. It consists in using a *tightening* operation before strengthening. The *tightening* operation is defined by:

Definition 3.14 [Tightening] Let B a DBM with elements in $\mathbb{Z}$. We define ${}^\sharp\mathcal{T}(B)$ be the DBM with elements in $\mathbb{Z}$ such that, for $u, v \in \mathbb{V}_\pm$:

$${}^\sharp\mathcal{T}(B)_{uv} = \begin{cases} B_{uv} - 1 & \text{if } u = \overline{v} \text{ and } B_{uv} \text{ is odd} \\ B_{uv} & \text{otherwise} \end{cases}$$

The following theorem gives the essential property of the tightening operation:

Theorem 3.15 *Let B a weakly closed DBM with elements in $\mathbb{Z}$. We suppose that $\forall u \in \mathbb{V}_\pm,\ 0 \leq {}^\sharp\mathcal{T}(B)_{u\overline{u}} + {}^\sharp\mathcal{T}(B)_{\overline{u}u}$. Then ${}^\sharp\mathcal{T}(B)$ is weakly tightly closed.*

Proof. See, e.g., [10, Theorem 8.2.18]. □

This theorem has two consequences. First, as already explained by Bagnara et al. [1, §6], it gives an efficient algorithm to compute tight closure: one would compute the closure of the input matrix, then tighten it and finally strengthen it. Second, our sparse algorithms need only small adjustments when used with integer environments: instead of maintaining the DBMs weakly closed, we just have to make them weakly tightly closed by tightening them after each operation.

Note, however, that tightening does not address the case of mixed environments, where some variables are known to have integer values, and some others can have an arbitrary real values. To the best of our knowledge, there is no known efficient closure algorithm supporting this use case, even in the dense setting.

4 Conclusion

In this paper, we presented new algorithms for the Octagon abstract domain, which preserve the sparsity of the representation of octagons. These algorithms are as

precise as the usual ones, and rely on a weaker invariant over difference bound matrices, called weak closedness. We have shown that these algorithms can be used in the context of rational or real environments as well as in the context of integer environments.

We implemented and formally verified in Coq these algorithms in the context of the Verasco static analyzer [9,10,18]. The use of these new algorithms improved the performances of the Octagon abstract domain by at least one order of magnitude.

There are still possible improvements to these algorithms: in particular, we think that it could be profitable to sparsify difference bound matrices as much as possible after each abstract operation, while still maintaining them weakly closed. Indeed, abstract operations may infer bounds in difference bound matrices that can actually be deduced from non-relational bounds, therefore missing opportunity of sparsity.

We think the reduction algorithm presented by Bagnara et al. [1] can be adapted to compute reduced difference bound matrices using only weakly closed difference bound matrices. This would lead to a simpler widening algorithm based on a semantic definition as described by Bagnara et al. [1, §4.2]. We believe the implementation of these new algorithms in state-of-the-art static analyzers, by using, for example, the framework developed by Singh et al. [17] would lead to a significant performance improvement.

References

[1] Bagnara, R., P. M. Hill and E. Zaffanella, *Weakly-relational shapes for numeric abstractions: Improved algorithms and proofs of correctness*, Formal Methods in System Design **35** (2009), pp. 279–323.

[2] Bouaziz, M., *TreeKs: A functor to make numerical abstract domains scalable*, in: *NSAD*, ENTCS **287** (2012), pp. 41–52.

[3] Chawdhary, A., E. Robbins and A. King, *Simple and efficient algorithms for octagons*, in: *APLAS*, LNCS **8858** (2014), pp. 296–313.

[4] Chen, L., J. Liu, A. Miné, D. Kapur and J. Wang, *An abstract domain to infer octagonal constraints with absolute value*, in: *SAS*, LNCS **8723** (2014), pp. 101–117.

[5] Clarisó, R. and J. Cortadella, *The octahedron abstract domain*, in: *SAS*, LNCS **3148** (2004), pp. 312–327.

[6] Cousot, P. and R. Cousot, *Static determination of dynamic properties of programs*, in: *Proceedings of the Second International Symposium on Programming* (1976), pp. 106–130.

[7] Cousot, P., R. Cousot, J. Feret, L. Mauborgne, A. Miné and X. Rival, *Why does astrée scale up?*, Formal Methods in System Design **35** (2009), pp. 229–264.

[8] Cousot, P. and N. Halbwachs, *Automatic discovery of linear restraints among variables of a program*, in: *POPL* (1978), pp. 84–96.

[9] Jourdan, J., V. Laporte, S. Blazy, X. Leroy and D. Pichardie, *A formally-verified C static analyzer*, in: *POPL*, ACM, 2015, pp. 247–259.

[10] Jourdan, J.-H., "Verasco: a Formally Verified C Static Analyzer," Ph.D. thesis, Université Paris Diderot (2016).

[11] Laviron, V. and F. Logozzo, *Subpolyhedra: A (more) scalable approach to infer linear inequalities*, in: *VMCAI*, LNCS **5403** (2009), pp. 229–244.

[12] Logozzo, F. and M. Fähndrich, *Pentagons: a weakly relational abstract domain for the efficient validation of array accesses*, in: *SAC*, ACM, 2008, pp. 184–188.

[13] Miné, A., "Weakly relational numerical abstract domains," Ph.D. thesis, École Polytechnique (2004).

[14] Miné, A., *The octagon abstract domain*, HOSC **19** (2006), pp. 31–100.

[15] Putot, S. and E. Goubault, *Static analysis of numerical algorithms*, in: *SAS*, LNCS **4134** (2006), pp. 18–34.

[16] Simon, A. and A. King, *The two variable per inequality abstract domain*, HOSC **23** (2010), pp. 87–143.

[17] Singh, G., M. Püschel and M. Vechev, *Making numerical program analysis fast*, in: *PLDI*, ACM, 2015, pp. 303–313.

[18] *The Verasco formally verified C static analyzer*, http://compcert.inria.fr/verasco/.